

09/29/00

09/29/00
JCS229 U.S. PTO
09/29/00

10-3-00

A

ASSISTANT COMMISSIONER FOR PATENTS
PATENT APPLICATION
Washington, D.C. 20231

CASE DOCKET NO. YOR920000464US1
Date: September 29, 2000

JCS229 U.S. PTO
09/29/00
09/29/00

Transmitted herewith for filing is the Patent Application of:

Inventors: Vaddekkadathu T. RAJAN et al.

For: MACHINE CUT TASK IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION

Enclosed are:

X 7 Sheets of Informal Drawings.

 An assignment of the invention to International Business Machines Corporation, Armonk, New York 10504.



 A certified copy of a application.


X Declaration and Power of Attorney is attached to the application.

 Associate Power of Attorney.

 Information Disclosure Statement with form PTO-1449 with references attached.

The filing fee has been calculated as shown below:

	(Col. 1)	(Col. 2)
FOR:	NO. FILED	NO. EXTRA
BASIC FEE		
TOTAL CLAIMS	27 - 20 =	7
INDEP CLAIMS	3 - 3 =	0
<u> </u> MULTIPLE DEPENDENT CLAIM PRESENTED		

OTHER THAN A SMALL ENTITY	
RATE	FEE
	\$ 690.00
X \$ 18 =	\$ 126.00
X \$ 78 =	\$ 0.00
+ \$ 260 =	\$ 0.00
TOTAL	\$ 816.00

If the difference in Col. 1 is less than zero, enter "0" in Col. 2.


x Please charge my Deposit Account No. 50-0510 in the amount of \$ 816.00.

x The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to Deposit Account No. 50-0510. A duplicate copy of this sheet is enclosed.

x Any additional filing fees required under 37 CFR 1.16.

x Any patent application processing fees under 35 CFR 1.17.

Respectfully submitted,

By 
Atty: Charles W. Peterson, Jr.
Registration No. 34,406
Tel. (202) 789-4900
Fax (202) 789-8707
FITCH, EVEN, TABIN & FLANNERY
120 South LaSalle Street
Chicago, Illinois 60603


Certificate of Mailing by Express Mail

"Express Mail" Mailing Label Number
EK754574885US

Date of Deposit 9-29-00

I hereby certify that this paper or fee is being Deposited with the United States Postal Service "Express Mail Post Office to Addressee" Service under 37 CFR Section 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, DC 20231

Charles W. Peterson, Jr.
(Typed or printed name of person mailing)


(Signature of person mailing)

MACHINE CUT TASK IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION

DESCRIPTION

RELATED APPLICATION

5 The present application is related to U.S. Patent Application No. 09/_____
(Attorney Docket No. YOR9-2000-0293-US1) entitled "INDEPENDENT NET TASK
IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION" to
Kimelman et al.; U.S. Patent Application No. 09/_____
10 (Attorney Docket No. YOR9-2000-0465-US1) entitled "NET ZEROING FOR EFFICIENT PARTITION AND
DISTRIBUTION" to Roth et al.; and U.S. Patent Application No. 09/_____
(Attorney Docket No. YOR9-2000-0466-US1) entitled "DOMINANT EDGE
IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION" to Wegman
et al. all filed coincident herewith and assigned to the assignee of the present invention.

BACKGROUND OF THE INVENTION

15

Field of the Invention

The present invention generally relates to distributed processing and more particularly, the present invention relates to efficiently assigning tasks across multiple computers for distributed processing.

Background Description

Any large, multifaceted project, such as a complex computer program, may be segmented into multiple smaller manageable tasks. The tasks then may be distributed amongst a group of individuals for independent completion, e.g., an engineering design project, distributed processing or, the layout of a complex electrical circuit such as a microprocessor. Ideally, the tasks are matched with the skills of the assigned individual and each task is completed with the same effort level as every other task. However, with such an ideal matched task assignment, intertask communication can become a bottleneck to project execution and completion. Thus, to minimize this potential bottleneck, it is important to cluster together individual tasks having the highest level of communication with each other. So, for example, in distributing eight equivalent tasks to pairs of individuals at four locations, (e.g., eight design engineers in four rooms) optimally, pairs of objects or tasks with the highest communication rate with each other are assigned to individual pairs at each of the four locations.

Many state of the art computer applications are, by nature, distributed applications. End-users sit at desktop workstations or employ palmtop information appliances on the run, while the data they need to access resides on distant data servers, perhaps separated from these end-users by a number of network tiers. Transaction processing applications manipulate data spread across multiple servers. Scheduling applications are run on a number of machines that are spread across the companies of a supply chain, etc.

When a large computer program is partitioned or segmented into modular components and the segmented components are distributed over two or more machines, for the above mentioned reasons, component placement can have a significant impact on

program performance. Therefore, efficiently managing distributed programs is a major challenge, especially when components are distributed over a network of remotely connected computers. Further, existing distributed processing management software is based on the assumption that the program installer can best decide how to partition the program and where to assign various program components. However, experience has shown that programmers often do a poor job of partitioning and component assignment.

So, a fundamental problem facing distributed application developers is application partitioning and component or object placement. Since communication cost may be the dominant factor constraining the performance of a distributed program, minimizing inter-system communication is one segmentation and placement objective. Especially when placement involves three or more machines, prior art placement solutions can quickly become unusable, i.e., what is known as NP-hard. Consequently, for technologies such as large application frameworks and code generators that are prevalent in object-oriented programming, programmers currently have little hope of determining effective object placement without some form of automated assistance. En masse inheritance from towering class hierarchies, and generation of expansive object structures leaves programmers with little chance of success in deciding on effective partitioning. This is particularly true since current placement decisions are based solely on the classes that are written to specialize the framework or to augment the generated application.

Furthermore, factors such as fine object granularity, the dynamic nature of object-based systems, object caching, object replication, ubiquitous availability of surrogate system objects on every machine, the use of factory and command patterns, etc., all make partitioning in an object-oriented domain even more difficult. In particular, for conventional graph-based approaches to partitioning distributed applications, fine-grained

object structuring leads to enormous graphs that may render these partitioning approaches impractical.

Finally, although there has been significant progress in developing middleware and in providing mechanisms that permit objects to inter-operate across language and machine boundaries, there continues to be little to help programmers decide
5 object-system placement. Using state of the art management systems, it is relatively straightforward for objects on one machine to invoke methods on objects on another machine as part of a distributed application. However, these state of the art systems provide no help in determining which objects should be placed on which machine in
10 order to achieve acceptable performance. Consequently, the initial performance of distributed object applications often is terribly disappointing. Improving on this initial placement performance is a difficult and time-consuming task.

Accordingly, there is a need for a way of automatically determining the optimal program segmentation and placement of distributed processing components to minimize
15 communication between participating distributed processing machines.

SUMMARY OF THE INVENTION

It is therefore a purpose of the present invention to improve distributed processing performance;

It is another purpose of the present invention to minimize communication between
20 distributed processing machines;

It is yet another purpose of the invention to improve object placement in distributed processing applications;

It is yet another purpose of the invention to determine automatically how objects should best be distributed in distributed processing applications

it is yet another purpose of the invention to minimize communication between objects distributed amongst multiple computers in distributed processing applications.

5 The present invention is a task management system, method and computer program product for determining optimal placement of task components on multiple machines for task execution, particularly for placing program components on multiple computers for distributed processing. First, a communication graph is generated representative of the computer program with each program unit (e.g., an object)
10 represented as a node in the graph. Nodes are connected to other nodes by edges representative of communication between connected nodes. A weight is applied to each edge, the weight being a measure of the level of communication between the connected edges. Terminal nodes representative of the multiple computers are attached to the communication graph. Independent nets may be separated out of the communication
15 graph. A cut is made at each terminal node and the weights of the cut edges are summed. The second heaviest terminal is identified from the cut and edges connected to at least one internal node and not connected to the second heaviest edge are compared against the weight of the second heaviest edge. Any edge found in the comparison to be at least as heavy as the second heaviest terminal node need not be included in the min cut for the
20 communication graph and so, is removed from consideration for the final min cut solution. Finally, program components which may be a single program unit or an aggregate of units are placed on computers according to the communication graph min cut solution.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed preferred embodiment description with reference to the drawings, in which:

5 Figure 1 shows an example of a flow diagram of the preferred embodiment of the present invention wherein a program is segmented, initially, and initial segments are distributed to and executed on multiple computers;

 Figures 2A-C show an example of a communication graph;

10 Figure 3 is a flow diagram of the optimization steps for determining an optimum distribution of program components;

 Figure 4 shows an example of a simple communication graph reducible by the preferred embodiment Machine Cut method of the present invention;

 Figure 5 is an example of the Machine Cut method steps of identifying non-terminal edges that may be removed from consideration;

15 Figure 6 is an example of the steps in contracting or collapsing edges that are at least as heavy as the second heaviest edge.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS OF THE INVENTION

20 As referred to herein, a communication graph is a graphical representation of a multifaceted task such as a computer program. Each facet of the task is an independent task or object that is represented as a node and communication between tasks or nodes is represented by a line (referred to as an edge) between respective communicating nodes. Participating individuals (individuals receiving and executing distributed tasks) are referred to as terminal nodes or machine nodes. A net is a collection of nodes connected

together by edges. Two nets are independent if none of the non-terminal nodes of one net shares an edge with a non-terminal node of the other. Thus, for example, a communication graph of a computer program might include a node for each program object and edges would be between communicating objects, with edges not being included between objects not communicating with each other. In addition, a weight indicative of the level of communication between the nodes may be assigned to each edge. Graphical representation and modeling of computer programs is well known in the art.

Referring now to the drawings, and more particularly, Figure 1 is an example of a flow diagram 100 of the preferred embodiment of the present invention wherein a program is segmented, initially, and initial segments are distributed to and executed on multiple computers. First, in step 102 the communication patterns of a program are analyzed to form a communication graph. Then, in step 104, the traces of the communication graph are analyzed, and an initial partition is determined. In step 106, the partition is optimized for minimum interpartition communication. In step 108, the individual objects are distributed amongst participants for execution according to the optimize partition of step 106.

A component refers to an independent unit of a running program that may be assigned to any participating individual, e.g., computers executing objects of a distributed program. Thus, a component may refer to an instance of an object in an object oriented program, a unit of a program such as Java Bean or Enterprise Java Bean or, a larger collection of program units or components that may be clustered together intentionally and placed on a single participating machine. Further, a program is segmented or partitioned into segments that each may be a single component or a collection of components. After segmenting, analyzing the segments and assigning each of segments

or components to one of the multiple participating machines or computers according to the present invention, the final machine assignment is the optimal assignment.

Thus, a typical communication graph includes multiple nodes representative of components with weighted edges between communicating nodes. Determining communication between components during program execution may be done using a typical well known tool available for such determination. Appropriate communication determination tools include, for example, Jinsight, for Java applications that run on a single JVM, the Object Level Tracing (OLT) tool, for WebSphere applications or, the monitoring tool in Visual Age Generator.

Figures 2A-C show an example of a communication graph of a net 110 that includes multiple nodes 112, 114, 116, 118, 120 and 122. Each node 112, 114, 116, 118, 120 and 122 represents a program component connected to communication edges 124, 126, 128, 130, 132, 134, 136 and 138 to form the net 110. Adjacent nodes are nodes that share a common edge, e.g., nodes 114 and 122 share edge 126. Each edge 124, 126, 128, 130, 132, 134, 136 and 138 has been assigned a weight proportional to, for example, the number of messages between the adjacent components.

In Figure 2B, Machine nodes 140, 142 and 144 representative of each participating machine (three in this example) are shown connected by edges 146, 148, 150. Initially, a node 112, 114, 116, 118, 120 and 122 may be placed on a machine 140, 142, 144 by adding an edge 146, 148, 150 with infinite weight (indicating constant communication) between the node and the machine. Typically, initial assignment places nodes with specific functions (e.g., database management) on a machine suited for that function. After the initial placement assigning some nodes 112, 114 and 122 to machines 140, 142, 144, other nodes 116, 118, 120 are assigned to machines 140, 142, 144, if they

communicate heavily with a node 112, 114, 122 already assigned to that machine 140, 142, 144. Additional assignment is effected by selectively collapsing edges, combining the nodes on either end of the collapsed edge and re-assigning edges that were attached to one of the two former adjacent nodes to the combined node. When assignment is
5 complete, all of the nodes 112, 114, 116, 118, 120 and 122 will have been placed on one of the machines at terminal nodes 140, 142, 144 and the final communication graph may be represented as terminal nodes 140, 142, 144 connected together by communication edges.

For this subsequent assignment, the graph is segmented by cutting edges and
10 assigning nodes to machines as represented by 152, 154 and 156 in Figure 2C to achieve what is known in the art as a minimal cut set or min cut set. A cut set is a set of edges that, if removed, eliminate every path between a pair of terminal nodes (machine nodes) in the graph. A min cut set is a cut set wherein the sum of the weights of the cut set edges is minimum. While there may be more than one min cut set, the sum is identical for all
15 min cut sets. A min cut may be represented as a line intersecting the edges of a min cut set. So, in the example of Fig. 2C, the sum of the weights of edges 124, 126, 128, 132 and 138 is 2090, which is cost of the cut and is representative of the total number of messages that would be sent between machines at terminal nodes 140, 142, 144 with this particular placement. The min cut identifies the optimum component placement with
20 respect to component communication. While selecting a min cut set may be relatively easy for this simple example, it is known to increase in difficulty exponentially with the number of nodes in the graph.

Figure 3 is a flow diagram 160 of the optimization steps for determining an optimum distribution of program components to individual participating computers
25 according to a preferred embodiment of the present invention. First, in step 162, an initial

communication graph is generated for the program. Then, in step 164 machine nodes are added to the communication graph. As noted above, certain types of components are designated, naturally, for specific host machine types, e.g., graphics components are designated for clients with graphics capability or, server components designated for a data base server. After assigning these host specific components, in step 168 independent nets are identified and the communication graph is partitioned into the identified independent nets as described in U.S. Patent Application No. 09/_____ (Attorney Docket No. YOR9-2000-0293-US1) entitled "INDEPENDENT NET TASK IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION" to Kimelman et al. assigned to the assignee of the present invention and incorporated herein by reference. In step 170 the Machine Cut reduction method described hereinbelow is used to reduce the independent nets and then, in step 172 a min cut for the reduced independent nets, the min cuts for all of the independent nets being the min cut for the whole communication graph.

Figure 4 example of a simple communication graph 180 reducible by the preferred embodiment Machine Cut method of the present invention. In this example, the graph 180 includes five (5) non-terminal nodes 182, 184, 186, 188 and 190 connected together by edges 192, 194, 196, 198, 200 and 202, referred to herein as non-terminal edges. Three (3) terminal nodes 204, 206 and 208 are connected to respective non-terminal nodes 182, 184, 186, 188 and 190 by edges 210, 212, 214, 216 and 218, referred to herein as non-terminal edges. A weight is represented as being attached to each edge 192 - 202 and 210 - 218. Dotted line 220 represents a terminal cut at terminal node 204 cutting terminal edges 210, 212. Dotted line 222 represents a terminal cut at terminal node 206 cutting terminal edges 214, 216. Dotted line 224 represents a terminal cut at terminal node 208 cutting terminal edge 218. Essentially, the Machine Cut method eliminates from inclusion in the min cut solution, any terminal or non-terminal edge with heavier communication (i.e., its weight exceeds) than all but the terminal node with the heaviest

level of communication. Thus, in this example, edge 202 is heavier than terminal cut 222. So, edge 202 can be excluded from consideration for inclusion in the min cut solution. Preferably, edge 202 is collapsed, combining nodes 182 and 188, as well as merging (then) parallel edges 198 and 200.

5 Figure 5 is an example of the Machine Cut method steps 230 of identifying non-terminal edges that may be removed from consideration according to the preferred embodiment of the present invention. First, in step 232, an independent net is selected for reduction. In step 234 terminal cuts are made at each terminal node on the selected net. For each terminal cut, the weights of the edges at the terminal are summed, the sum
10 being the terminal's weight. Then, in step 236, the second heaviest terminal node (the terminal with the second heaviest weight) is identified. In step 238, edges are checked to determine if they are at least as heavy as the identified second heaviest cut weight. All edges connected to at least one non-terminal node are checked in step 238, except that those edges connected to the second heaviest node are excluded. If no edges are found
15 that are as heavy or heavier than the second heaviest cut weight, in step 240, it is determined that the Machine Cut method is unable to reduce the net and in step 242, net reduction ends. Otherwise, in step 244, each edge that was identified in step 238 need not be part of the min cut solution and so, is collapsed. In step 246 it is determined that the independent net has been reduced using the Machine Cut method 230 and, net reduction
20 ends in step 242.

 In other words, for each terminal node 204, 206, 208 the weight of all connected edges are summed. Then, the summed are sorted in descending order and the second largest weight is selected and labeled W_{2nd} , for example. Next, any edge 192 - 202, 210, 212 and 218 not connected to the second heaviest node but connected to at least one
25 non-terminal node are compared against W_{2nd} . Any compared edge that is at least as

heavy as W_{2nd} need not be part of the (only) multiway minimum cut of the graph, and as such, may be collapsed. Collapsing each edge results in a simpler graph wherein the min cut solution can be found much more quickly and efficiently, with the min cut solution weight being the same as the original unreduced graph.

5 Figure 6 is an example of the steps in contracting or collapsing edges that are at least as heavy as the second heaviest edge 244. First in step 2440 the two nodes connected by the collapsed edge are merged, resulting in a single merged node that includes the components of both original nodes. Then, in step 2442 the collapsed edge is discarded. Finally, in step 2444 any “parallel” edge groups (edges connecting the merged node to the same adjacent node) resulting from the merger are replaced with a single edge with its weight equal to the sum of parallel edge weights. Thus, as a result of contracting dominant edges, the graph has been reduced wherein a min cut solution may be found with less effort.

10 In the preferred embodiment, the min cut step 170 is an iterative process, wherein independent nets are reduced using the Machine Cut steps described herein and, when necessary, in combination with other linear complexity methods such as the Dominant Edge identification method of U.S. Patent Application No. 09/_____ (Attorney Docket No. YOR9-2000-0466-US1) entitled “DOMINANT EDGE IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION” to Wegman et al. and the Net Zeroing method of U.S. Patent Application No. 09/_____ (Attorney Docket No. YOR9-2000-0465-US1) entitled “NET ZEROING FOR EFFICIENT PARTITION AND DISTRIBUTION” to Roth et al., all filed coincident herewith, assigned to the assignee of the present invention and incorporated herein by reference. Further, as independent nets are reduced, those reduced nets are further checked as in step 168 above to determine if they may be divided into simpler independent nets. Then, the Machine Cut method of the

preferred embodiment is applied to those simpler independent nets. To reach a solution more quickly, on each subsequent pass, only nodes and edges of a subgraph that were adjacent to areas reduced previously are rechecked. Thus, the communication graph is simplified by eliminating machine cut edges to reach a min cut solution much quicker and much more efficiently than with prior art methods.

The reduction method of the preferred embodiment reduces the number of independent components in the communication graph of a complex program. In the best case, an appropriate allocation of every component in the program is provided. However, even when best case is not achieved, the preferred embodiment method may be combined with other algorithms and heuristics such as the branch and bound algorithm or the Kernighan-Lin heuristic to significantly enhance program performance. Experimentally, the present invention has been applied to communication graphs of components in several programs with results that show significant program allocation improvement, both in the quality of the final solution obtained and in the speed in reaching the result.

Although the preferred embodiments are described hereinabove with respect to distributed processing, it is intended that the present invention may be applied to any multi-task project without departing from the spirit or scope of the invention. Thus, for example, the task partitioning and distribution method of the present invention may be applied to VLSI design layout and floor planning, network reliability determination, web pages information relationship identification, and "divide and conquer" combinatorial problem solution approaches, e.g., "the Traveling Salesman Problem."

While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

CLAIMS

What is claimed is:

1. A task management method for determining optimal placement of task components, said method comprising:

5 a) generating a communication graph representative of a task, task components represented as nodes of said communication graph and edges connecting ones of said nodes, said edges representing communication between connected nodes and being weighted proportional to communication between connected nodes;

b) assigning terminal nodes to said communication graph;

10 c) identifying high communication edges within said communication graph, said high communication edges having a weight indicating a communication level exceeding the communication level for a selected terminal node;

d) determining a min cut solution for said communication graph, high communication edges being excluded from determined min cut solutions; and

15 e) placing task components on said terminal nodes responsive to said min cut solution.

2. A task management method as in claim 1, after the step (b) of assigning terminal nodes, further comprising the step of:

20 b1) identifying independent nets in said communication graph, each of said independent nets being connected between a plurality of said terminal nodes.

3. A task management method as in claim 2, wherein the step (c) of identifying high communication edges comprises the steps of:

i) summing the weight of terminal edges connected to terminal nodes of an independent net;

5 ii) identifying the terminal node having the second largest sum as the second heaviest terminal node;

iii) identifying any edge connected to at least one non-terminal node and not connected to said second heaviest node and at least as heavy as the second largest sum; and

10 iv) collapsing each identified edge.

4. A task management method as in claim 3, further comprising the step of:

v) repeating steps (i) - (iv) until no edges are identified as being heavier than the second largest sum.

15 5. A task management method as in claim 4, wherein identified edges are selectively collapsed comprising the steps of:

i) merging nodes at opposite end of each identified edge to form a single merged node including the components of both original nodes;

ii) discarding the identified edge; and

20 iii) replacing groups of parallel edges with a single edge having a weight equal to the sum of parallel edge weights.

6. A task management method as in claim 5, wherein the step (c) of determining a min cut solution comprises the steps of:

i) identifying independent nets in reduced nets;

ii) identifying and collapsing edges selectively identified as being heavier than the second heaviest terminal node in said identified independent nets, said independent nets being further reduced; and

iii) repeating steps (i) - (ii) until a min cut solution has been found.

5 7. A task management method as in claim 6, wherein each said task component is a unit of the computer program.

8. A task management method as in claim 7, wherein said each computer program unit is an instance of an object in an object oriented program

10 9. A task management method as in claim 7, wherein in step (d) computer program units are placed on computers, computer program units being placed on a common computer being combined into a single component.

10. A task management method as in claim 6 wherein said task is integrated circuit chip functional element placement and said task components are logic elements, said logic elements being placed on an integrated circuit chip in placement step (e).

15 11. A distributed processing system for determining optimal placement of computer program components on multiple computers, said distributed processing system comprising:

means for generating a communication graph of nodes interconnected by edges and representative of a computer program, computers executing said computer program
20 being represented as terminal nodes, computer program components being represented as non-terminal nodes, said edges representing communication between connected nodes and being weighted proportional to communication between connected nodes;

means for summing the weight of edges connected to terminal nodes;
means for identifying a second heaviest terminal node;
means for comparing edges with the sum for said second heaviest terminal node;
means for determining a min cut solution for said communication graph, edges

5 heavier than said sum being excluded from determined min cut solutions responsive to said comparison; and

means for placing program components on ones of said computers responsive to said determined min cut solution; and

said computer program being executed by said computers.

10 12. A distributed processing system as in claim 11, further comprising:
means for identifying independent nets connected between a plurality of said terminal nodes.

13. A distributed processing system as in claim 12, further comprising:
means for collapsing said edges heavier than said sum.

15 14. A distributed processing system as in claim 13, wherein the means for identifying edges heavier than said sum comprises:

means for summing the weight of terminal edges connected to terminal nodes;
means for identifying the terminal node having the second largest sum as the second heaviest terminal node;

20 means for comparing edge weights against said second largest sum; and
means for selectively collapsing edges identified as having a weight at least as heavy as the second largest sum.

15. A distributed processing system as in claim 14, the means for selectively collapsing edges further comprising:

means for merging nodes on either end of a selected edge and discarding said selected edge; and

means for replacing pairs of parallel edges attached to said merged node with a single edge.

16. A distributed processing system as in claim 15, wherein the means for comparing edge weights further comprises:

means for selecting edges attached to at least one non-terminal node and not attached to said second heaviest terminal node.

17. A distributed processing system as in claim 16, wherein each said program component is a unit of the computer program.

18. A distributed processing system as in claim 17, wherein said each program unit is an instance of an object in an object oriented program

19. A computer program product for determining optimal placement of functional components, said computer program product comprising a computer usable medium having computer readable program code thereon, said computer readable program code comprising:

computer readable program code means for generating a communication graph of nodes interconnected by edges and representative of a function, a plurality of said nodes being terminal nodes, functional components being represented as non-terminal nodes, said edges representing communication between connected nodes and being weighted proportional to communication between connected nodes;

computer readable program code means for summing the weight of edges connected to terminal nodes;

computer readable program code means for identifying a second heaviest terminal node;

5 computer readable program code means for comparing edges with the sum for said second heaviest terminal node;

computer readable program code means for determining a min cut solution for said communication graph, edges heavier than said second heaviest edge being excluded from determined min cut solutions responsive to said comparison; and

10 computer readable program code means for placing functional components responsive to said determined min cut solution.

20. A computer program product as in claim 19, further comprising:

computer readable program code means for identifying independent nets connected between a plurality of said terminal nodes.

15 21. A computer program product as in claim 20, further comprising:

computer readable program code means for collapsing edges heavier than said sum.

22. A computer program product as in claim 21, wherein the computer readable program code means for identifying edges heavier than said sum comprises:

20 computer readable program code means for summing the weight of terminal edges connected to terminal nodes;

computer readable program code means for identifying the terminal node having the second largest sum as the second heaviest terminal node;

computer readable program code means for comparing edge weights against said second largest sum; and

computer readable program code means for selectively collapsing edges identified as having a weight at least as heavy as the second largest sum.

5 23. A computer program product as in claim 22, wherein the computer readable program code means for selectively collapsing edges further comprising:

 computer readable program code means for merging nodes on either end of a selected edge and discarding said selected edge; and

10 computer readable program code means for replacing pairs of parallel edges attached to said merged node with a single edge.

24. A computer program product as in claim 23, wherein the computer readable program code means for comparing edge weights further comprises:

 computer readable program code means for selecting edges attached to at least one non-terminal node and not attached to said second heaviest terminal node.

15 25. A computer program product as in claim 24, wherein said function is a computer program and each said functional component is a unit of the computer program.

26. A computer program product as in claim 25, wherein each said program unit is an instance of an object in an object oriented program.

20 27. A computer program product as in claim 24 wherein said function is an integrated circuit chip and said functional components are logic elements.

MACHINE CUT TASK IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION

ABSTRACT OF THE INVENTION

5 A task management system, method and computer program product for
determining optimal placement of task components on multiple machines for task
execution, particularly for placing program components on multiple computers for
distributed processing. First, a communication graph is generated representative of the
computer program with each program unit (e.g., an object) represented as a node in the
graph. Nodes are connected to other nodes by edges representative of communication
10 between connected nodes. A weight is applied to each edge, the weight being a measure
of the level of communication between the connected edges. Terminal nodes
representative of the multiple computers are attached to the communication graph.
Independent nets may be separated out of the communication graph. A cut is made at
each terminal node and the weights of the cut edges are summed. The second heaviest
15 terminal is identified from the cut and edges connected to at least one internal node and
not connected to the second heaviest edge are compared against the weight of the second
heaviest edge. Any edge found in the comparison to be at least as heavy as the second
heaviest terminal node need not be included in the min cut for the communication graph
and so, is removed from consideration for the final min cut solution. Finally, program
20 components which may be a single program unit or an aggregate of units are placed on
computers according to the communication graph min cut solution.

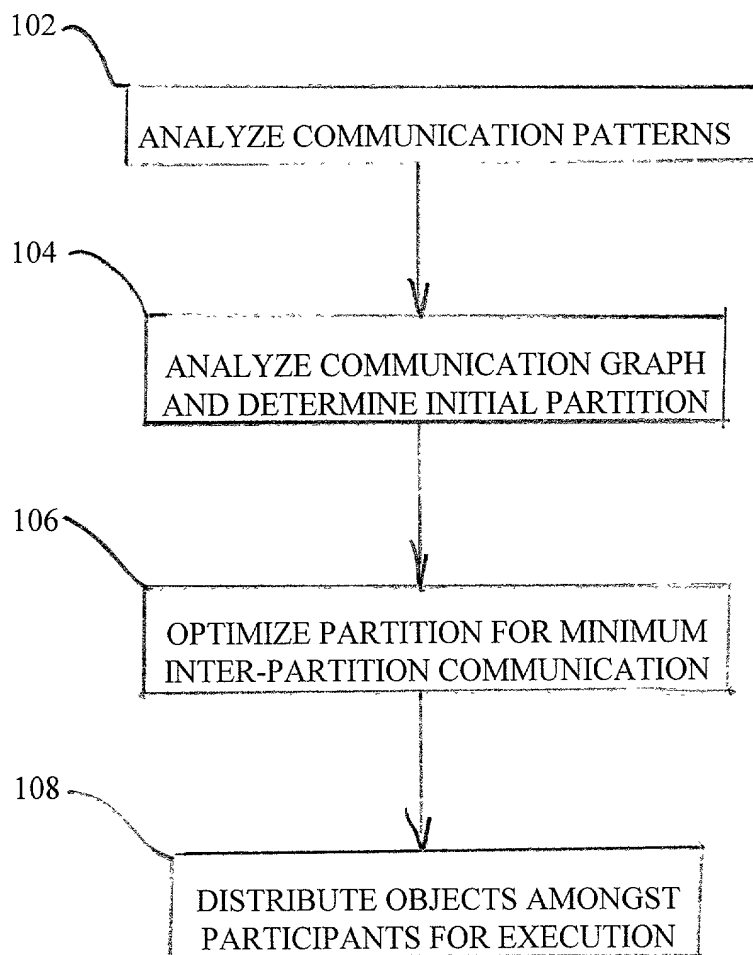




FIG. 1



006250" E2492950




Fig. 2C

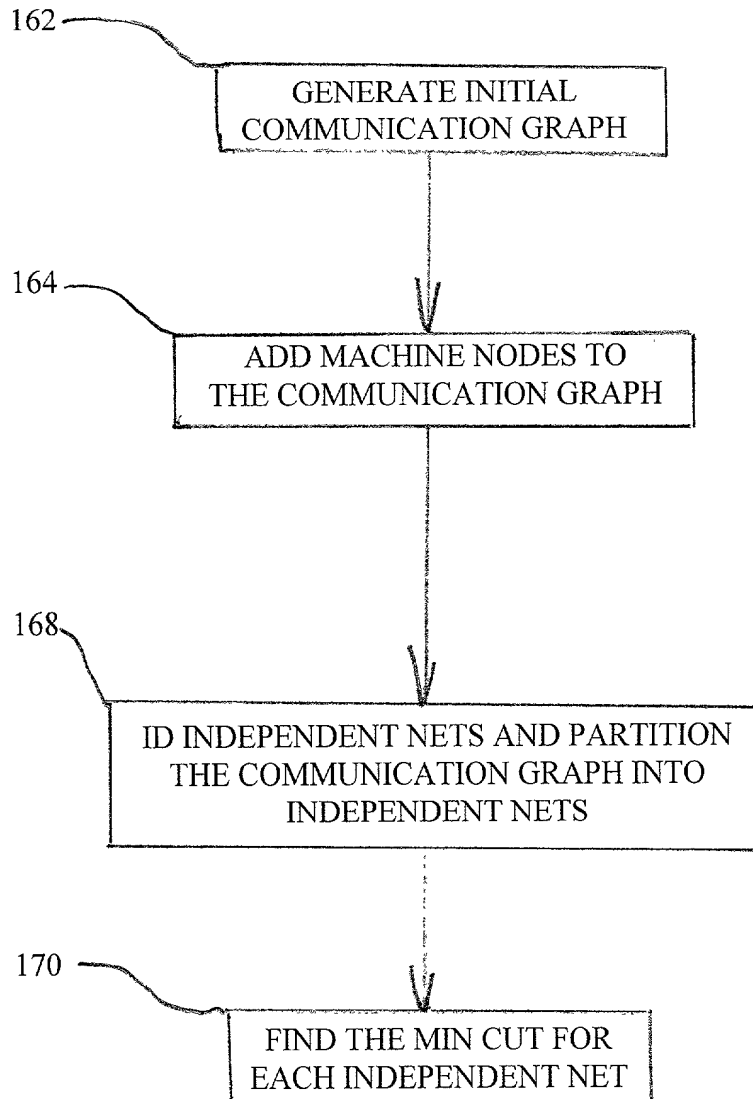


FIG. 3


$$W_2 = 30$$

$$W_1 = 40$$

Fig. 4




Fig. 6

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

MACHINE CUT TASK IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION

the specification of which (check one)

☒ X is attached hereto.

_____ was filed on _____ as United States Application Number

or PCT International Application Number _____

and was amended on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119(a)-(d) or §365(b) of any foreign application(s) for patent or inventor's certificate, or §365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application, having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

Priority Claimed

(Number)

(Country)

(Day/Month/Year Filed)

☐ Yes ☐ No

(Number)

(Country)

(Day/Month/Year Filed)

☐ Yes ☐ No

(Number)

(Country)

(Day/Month/Year Filed)

☐ Yes ☐ No

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below.

(Application Number)

(Filing Date)

(Application Number)

(Filing Date)

I hereby claim the benefit under 35 U.S.C. §120 of any United States Application(s), or §365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States, or PCT International application in the manner provided by the first paragraph of 35 U.S.C. §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in 37 CFR §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial No.)

(Filing Date)

(Status) (patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status) (patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith (list name and registration number).

Manny W. Schecter (Reg. 31,722), Lauren C. Bruzzzone (Reg. 35,082), Christopher A. Hughes (Reg. 26,914), John E. Hoel (Reg. 26,279), Joseph C. Redmond, Jr. (Reg. 18,753), Stephen C. Kaufman (Reg. 29,551), Jay P. Sbrrollini (Reg. 36,266), David M. Shofi (Reg. 39,835), Robert M. Trepp (Reg. 25,993), Louis P. Herzberg (Reg. 41,500), Daniel P. Morris (Reg. 32,053), Paul J. Otterstedt (Reg. 37,411), Louis J. Percello (Reg. 33,206), Douglas W. Cameron (Reg. 31,596), Wayne L. Ellenbogen (Reg. 43,602) and Marian Underweiser (Reg. 46,134)

Send Correspondence to: Charles W. Peterson, Jr. Fitch, Even, Tabin & Flannery120 South LaSalle Street, Suite 1600, Chicago, IL 60603Direct Telephone Calls to: (name and telephone number) Charles W. Peterson, (202) 789-4900Vadekkadathu T. RAJAN

Full name of sole or first inventor

Inventor's Signature

Date

223 Schrade Road, Briarcliff Manor, NY 10510

Residence

USA

Citizenship

Same as above

Post Office Address

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATIONDouglas N. KIMELMAN

Full name of second joint-inventor, if any

Inventor's signature

Date

P.O. Box 704, Yorktown Heights, NY 10510
ResidenceCanada
Citizenshipsame as above
Post Office AddressTova ROTH

Full name of third joint-inventor, if any

Inventor's signature

Date

408 Edward Avenue, Woodmere, NY 11598
ResidenceUSA
CitizenshipSame as above
Post Office AddressMark N. WEGMAN

Full name of fourth joint-inventor, if any

Inventor's Signature

Date

43 Overlook Road, Ossining, NY 10562
ResidenceUSA
CitizenshipSame as above
Post Office AddressKarin HOGSTEDT

Full name of fifth joint inventor, if any

Inventor's Signature

Date

Room D275, 180 Park Avenue, Florham Park, NJ 07932
ResidenceSweden
CitizenshipSame as above
Post Office Address

Full name of sixth joint-inventor, if any

Inventor's signature

Date

Residence

Citizenship

Post Office Address